

Scott Pakin

Research interests

Computer architecture, with an emphasis on memory subsystems; high-speed communication networks; and, parallel computer systems

Education

Ph.D. October, 2001 U. of Illinois at Urbana-Champaign GPA: 3.81/4.00
Comp. Sci. *Thesis title:* “Unresponsiveness-tolerant collective communication”

M.S. January, 1995 U. of Illinois at Urbana-Champaign GPA: 3.81/4.00
Comp. Sci. *Thesis title:* “The impact of message traffic on multicomputer memory hierarchy performance”

B.S. May, 1992 Carnegie Mellon University GPA: 3.42/4.00
Math/CS

Experience

Los Alamos National Laboratory Adolfo Hoisie (505) 667-5216 2/02–present

- ★ Created and implemented a novel, domain-specific programming language designed specifically for benchmarking and validating communication networks.
- ★ Developed a new technique for analyzing the impact of the memory wall on large-scale applications.
- ★ Assisted in the benchmarking and analysis of various production and preproduction parallel computer systems.

NCSA Rob Pennington (217) 244-1052 8/01–2/02

- ★ Collaborated on a project to develop NCSA’s next-generation communication middleware, the Virtual Machine Interface 2.0. When completed, VMI 2.0 will run on all of the production clusters at the National Center for Supercomputing Applications.

U. of Illinois at Urbana-Champaign Andrew Chien (858) 822-2458 1/95–8/01

- ★ Led development team producing Illinois Fast Messages (FM), a fast, efficient messaging layer that provides the communication infrastructure for the National Center for Supercomputing Applications’ NT Supercluster, at the time one of the world’s 500 fastest computing platforms. FM was also one of four academic research projects influencing Compaq/Intel/Microsoft’s Virtual Interface Architecture (VIA) standard for cluster interconnects.

(Continued on next page)

- ★ Cooperated with researchers at MIT on implementing a novel coordinated process scheduling algorithm that improves the throughput and response time of parallel and distributed applications running on clusters of timeshared workstations.
- ★ Designed and implemented a new technique for improving the performance of collective-communication-intensive parallel applications running on timeshared PC clusters.
- ★ Supervised various undergraduate research projects.

Intel Corporation **Tom Shott** **(503) 677-5019** **5/95–8/95**

- ★ Helped architect and validate a commercial, programmable, cache-coherent, scalable distributed shared memory system based on Intel Pentium Pro processors.
- ★ Defined and implemented a tool to ensure the the behavioral, functional, and register-transfer level simulators correctly model the hardware, software, and firmware used to implement the system's cache-coherency protocol.
- ★ Assisted with development and refinement of the protocol.

Awards and Honors

Named one of the top 5 new hires of the 2001–2002 fiscal year by a Los Alamos National Laboratory Director's committee

Honor Society of Phi Kappa Phi for superior academic performance. Inducted 1999.

Intel Foundation Graduate Fellowship for doctoral students at selected universities who are judged to be the “best of the best” in fields of study related to Intel's technology research. September, 1998–May, 1999.

W. J. Poppelbaum Award for the CS department's most outstanding graduate student in the areas of hardware and architecture. March, 1997.

Graduate College Fellowship tuition and fee waiver for the best three CS graduate students not yet on an assistantship. September, 1993–May, 1994.

(Continued on next page)

Publications

Journal

Andrew A. Chien, Mario Lauria, Rob Pennington, Mike Showerman, Giulio Iannello, Matt Buchanan, Kay Connelly, Louis Giannini, Greg Koenig, Sudha Krishnamurthy, Qian Liu, Scott Pakin, and Geetanjali Sampemane. Design and evaluation of an HPVM-based Windows NT supercomputer. *International Journal of High Performance Computing Applications*, 13(3):201–219, Fall 1999. Special issue on clusters and computational grids for scientific computing. Available from <http://www.c3.lanl.gov/~pakin/papers/bbfarm.ps>.

Mario Lauria, Scott Pakin, and Andrew Chien. Efficient layering for high speed communication: The MPI over Fast Messages (FM) experience. *Cluster Computing*, 2:107–116, 1999. Available from <http://www.c3.lanl.gov/~pakin/papers/cluster.ps>.

Scott Pakin, Vijay Karamcheti, and Andrew A. Chien. Fast Messages: Efficient, portable communication for workstation clusters and MPPs. *IEEE Concurrency*, 5(2):60–73, April–June 1997. Available from <http://ieeexplore.ieee.org/iel1/4434/12910/00588295.pdf>. (A draft is available from <http://www.c3.lanl.gov/~pakin/papers/fm-pdt.ps>).

Conference and Workshop

Eitan Frachtenberg, Fabrizio Petrini, Juan Fernandez, Scott Pakin, and Salvador Coll. STORM: Lightning-fast resource management. In *Proceedings of the IEEE/ACM SC2002 Conference*, Baltimore, Maryland, November 16–22, 2002. Available from <http://www.c3.lanl.gov/~pakin/papers/sc2002.pdf>.

Darren Kerbyson, Adolfo Hoisie, Scott Pakin, Fabrizio Petrini, and Harvey Wasserman. Performance testing of an EV7 AlphaServer machine. In *Proceedings of the Los Alamos Computer Science Institute (LACSI) Symposium*, Santa Fe, New Mexico, October 13–16, 2002. Available from http://www.c3.lanl.gov/par_arch/pubs/LAUR-02-4850.pdf.

Scott Pakin and Avneesh Pant. VMI 2.0: A dynamically reconfigurable messaging layer for availability, usability, and management. In *The 8th International Symposium on High Performance Computer Architecture (HPCA-8), Workshop on Novel Uses of System Area Networks (SAN-1)*, Cambridge, Massachusetts, February 2, 2002. Available from <http://www.csl.cornell.edu/SAN-1/san1.pdf>.

Geetanjali Sampemane, Scott Pakin, and Andrew A. Chien. Performance monitoring on an HPVM cluster. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2000), International Workshop on Cluster Computing—Technologies, Environments, and Applications (CC-TEA 2000)*, Las Vegas, Nevada, June 26–29, 2000. Available from <http://www.ac.upc.es/homes/toni/CCTEA2000/Sampemane.pdf>.

Mario Lauria, Scott Pakin, and Andrew A. Chien. Efficient layering for high speed com-

(Continued on next page)

munication: Fast Messages 2.x. In *Proceedings of the Seventh IEEE International Symposium on High Performance Distributed Computing (HPDC-7)*, pages 10–20, Chicago, Illinois, July 28–31, 1998. Available from <http://www.c3.lanl.gov/~pakin/papers/hpdc7-lauria.ps>.

Patrick Sobalvarro, Scott Pakin, Andrew Chien, and William Weihl. Dynamic coscheduling on workstation clusters. In *12th Annual International Parallel Processing Symposium & 9th Symposium on Parallel and Distributed Processing (IPPS/SPDP)*, 4th Workshop on Job Scheduling Strategies for Parallel Processing, Orlando, Florida, March 1998. Published in Lecture Notes in Computer Science, vol. 1459, pp. 231–256. Springer-Verlag. ISBN 3-540-64825-9. Available from <http://link.springer.de/link/service/series/0558/papers/1459/14590231.pdf>.

Andrew Chien, Scott Pakin, Mario Lauria, Matt Buchanan, Kay Hane, Louis Giannini, and Jane Prusakova. High performance virtual machines (HPVM): Clusters with supercomputing APIs and performance. In Michael Heath, Virginia Torczon, Greg Astfalk, Petter E. Bjørstad, Alan H. Karp, Charles H. Koebel, Vipin Kumar, Robert F. Lucas, Layne T. Watson, and David E. Womble, editors, *Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing*, Minneapolis, Minnesota, March 1997. Available from <http://www.c3.lanl.gov/~pakin/papers/hpvm-siam97.ps>.

Scott Pakin, Mario Lauria, and Andrew Chien. High performance messaging on workstations: Illinois Fast Messages (FM) for Myrinet. In *Proceedings of the 1995 ACM/IEEE Supercomputing Conference*, volume 2, pages 1528–1557, San Diego, California, December 1995. Available from http://www.supercomp.org/sc95/proceedings/567_SPAK/SC95.PDF.

Other

Patrick Sobalvarro, Scott Pakin, Andrew Chien, and William Weihl. Dynamic coscheduling on workstation clusters. SRC Technical Note 1997-017, Digital Equipment Corporation, Systems Research Center, 130 Lytton Avenue, Palo Alto, California 94301, March 14, 1997. Available from <ftp://gatekeeper.research.compaq.com/pub/DEC/SRC/technical-notes/SRC-1997-017.pdf>.

Scott Pakin. Regular expressions and gender guessing. *Computer Language Magazine*, 8(12):59–68, December 1991.

Public Talks

“Performance evaluation of a GS1280 ‘Marvel’ prototype”, The Consortium for Advanced Scientific and Technical Computing Users Group (CAST) Spring 2003 Meeting (March, 2003)

Invited talk, “High-performance Computing on Commodity-based Clusters”, Department of Computer and Information Science, The Ohio State University (February, 2002)

(Continued on next page)

“High Performance Computing on Windows NT” Birds of a Feather session at SC ’97 (November, 1997)—moderated and presented

“Systems management tools/methodology” talk and panel discussion in Sandia National Laboratories’ first Scalable Cluster Workshop (November, 1997)

“High performance virtual machines (HPVM): Clusters with supercomputing APIs and performance” paper presentation at the Eighth SIAM Conference on Parallel Processing for Scientific Computing (March, 1997)

“Works in Progress/Outrageous Opinions” session at the ACM Seventh International Conference on Architectural Support for Programming Languages and Operating Systems (October, 1996)

“Communications Layer ‘Shootout’” panel at the ACM Seventh International Conference on Architectural Support for Programming Languages and Operating Systems’ NOW/cluster workshop (October, 1996)

Intel Research Council University Forum, “Stepping Off the Bus: Clusters of Servers & Networks of Workstations” (June, 1996)

“High performance messaging on workstations: Illinois Fast Messages (FM) for Myrinet” paper presentation at the 1995 ACM/IEEE Supercomputing Conference (December, 1995)

Professional Activities

Program committee member, The Twelfth IEEE International Symposium on High Performance Distributed Computing (HPDC-12), 2003

Program committee member, Communication Architecture for Clusters (CAC) workshop, The 2003 International Parallel and Distributed Processing Symposium (IPDPS 2003)

Session chair, “Communication Libraries”, Communication Architecture for Clusters (CAC) workshop, The 2003 International Parallel and Distributed Processing Symposium (IPDPS 2003)

Session chair, “Routing and Switching”, Communication Architecture for Clusters (CAC) workshop, The 2002 International Parallel and Distributed Processing Symposium (IPDPS 2002)

Program committee member, Communication Architecture for Clusters (CAC) workshop, The 2002 International Parallel and Distributed Processing Symposium (IPDPS 2002)

Other

U.S. citizen; active U.S. Department of Energy Q-level security clearance